Bayesian Auctions with Efficient Queries (Extended Abstract)*

Jing Chen[1], **Bo Li**[2], **Yingkai Li**[3] and **Pinyan Lu**[4]
[1]Department of Computer Science, Stony Brook University
[2]Department of Computing, The Hong Kong Polytechnic University
[3]Department of Computer Science, Northwestern University
[4]Institute for Theoretical Computer Science, Shanghai University of Finance and Economics
jingchen@cs.stonybrook.edu, comp-bo.li@polyu.edu.hk, yingkai.li@u.northwestern.edu,
lu.pinyan@mail.shufe.edu.cn

Abstract

Designing dominant-strategy incentive compatible (DSIC) mechanisms for a seller to generate optimal revenue by selling items to players is a fundamental problem in Bayesian mechanism design. However, most existing studies assume that the seller knows the entire distribution from which the players' values are drawn, which may not hold in reality. In this work we consider, for the first time, the *query complexity* of Bayesian mechanisms. The seller has limited oracle accesses to the players' distributions, via *quantile queries* and *value queries*. For single-item auctions, we design mechanisms with *logarithmic* number of value or quantile queries which achieve almost optimal revenue. We then prove logarithmic lower-bounds, i.e., logarithmic number of queries are necessary for any constant approximation DSIC mechanisms, even when randomized and adaptive queries are allowed. Thus our mechanisms are almost optimal regarding query complexity. Our lower-bounds can be extended to multi-item auctions with monotone subadditive valuations, and we complement this part with constant approximation mechanisms for unit-demand or additive valuation functions. Our results are robust even if the answers to the queries contain noises.

1 Introduction

An important problem in Bayesian mechanism design is to design auctions that (approximately) maximize the seller's expected revenue by selling items. More precisely, in a Bayesian multi-item auction a seller has m heterogeneous items to sell to n players. Each player has a private value for each item, which is independently drawn from some prior distribution. When the joint prior distribution is of *common knowledge* to both the seller and the players, optimal Bayesian incentive-compatible (BIC) mechanisms have been discovered for various auction settings [Myerson, 1981; Cremer and McLean, 1988; Cai *et al.*, 2012], where all players reporting their true values forms a Bayesian Nash equilibrium. When there is no common prior but the seller knows the distribution, many (approximately) optimal dominant-strategy incentive-compatible (DSIC) Bayesian mechanisms have been designed [Myerson, 1981; Ronen, 2001; Chawla *et al.*, 2010; Cai *et al.*, 2016], where it is each player's *dominant strategy* to report his true values.

However, the *complexity* for the seller to carry out such mechanisms is largely unconsidered in the literature. Most existing Bayesian mechanisms require that the seller has full access to the prior distribution and is able to carry out all required optimizations based on the distribution, so as to compute the allocation and the prices. Unfortunately the seller may not be so knowledgeable or powerful in real-world scenarios. If the supports of the distributions are exponentially large (in m and n), or if the distributions are continuous and do not have succinct representations, it is hard for the seller to write out "each single bit" of the distributions or precisely carry out arbitrary optimization tasks based on them. Even with a single player and a single item, when the value distribution is irregular, computing the optimal price in time that is much smaller than the size of the support is not an easy task. Thus, a natural and important question to ask is *how much the seller should know about the distributions in order to obtain approximately optimal revenue.*

In this work we consider, for the first time, the *query complexity* of Bayesian mechanisms. In particular, the seller can only access the distributions by making oracle queries. Two types of queries are allowed, *quantile queries* and *value queries*. That is, the seller queries the oracle with specific quantiles (respectively, values), and the oracle returns the corresponding values (respectively, quantiles) in the underlying distributions. These two types of queries happen a lot in market study. Indeed, the seller may wish to know what is the price he should set so that half of the consumers would purchase his product; or if he sets the price to be 200 dollars, how many consumers would buy it. Another important scenario where such queries naturally come up is in databases. Indeed, although the seller may not know the distribution, some powerful institutes, say the Office for National Statistics, may have such information figured out and stored in its database. As in most database applications, it may be neither necessary nor feasible for the seller to download the whole distribution

*The full paper is published in Artificial Intelligence 303 (2022) 103630. A brief announcement has appeared on 45th International Colloquium on Automata, Languages, and Programming.

to his local machines. Rather, he would like to access the distribution via queries to the database. Moreover, the commercial data providers may charge the mechanism designer a fixed monetary payment for each query about the statistics of the distribution. To maximize the revenue of the designer, it is necessary to understand the tight bounds on the number of queries that are sufficient for good revenue guarantees.

Another concern in practice is that the queries to the valuation distributions are not precise. It is possible that the data provider only has estimates for the statistics of the valuation distributions. and those estimates are only correct with high probability. The data provider may even want to obfuscate the data ue to privacy issues [Dwork, 2008] We show that our results are robust with the presence of imprecise queries.

2 Preliminaries

2.1 Bayesian Auctions

In a multi-item auction there are m items, denoted by $M = \{1, \ldots, m\}$, and n players, denoted by $N = \{1, \ldots, n\}$. Each player $i \in N$ has a non-negative value for each item $j \in M$, v_{ij}, which is independently drawn from distribution D_{ij}. Player i's *true valuation* is $(v_{ij})_{j\in[m]}$. To simplify the notations, we may write v_i for $(v_{ij})_{j\in[m]}$ and v for $(v_i)_{i\in[n]}$. Letting $D_i = \times_{j\in M} D_{ij}$ and $D = \times_{i\in N} D_i$, we use $\mathcal{I} = (N, M, D)$ to denote the corresponding Bayesian auction instance. We will consider several classes of widely studied auctions. A *single-item* auction has $m = 1$. When $m > 1$, a player i being *unit-demand* means his value for a subset S of items is $\max_{j\in S} v_{ij}$, and a player i being *additive* means his value for S is $\sum_{j\in S} v_{ij}$. When all players are unit-demand (respectively, additive), we call such an auction a *unit-demand auction* (respectively, an *additive auction*).

A mechanism $\mathcal{M}$ maps a reported value profile v from the players to a/n (random) allocation of items, $x = (x_i(v))_{i\in N}$, and payments to charge the players, $p = (p_i(v))_{i\in N}$. When v is clear from the context, we simply denote by $x = (x_i)_{i\in N}$ and $p = (p_i)_{i\in N}$. For single-item auctions, $x_i \in [0, 1]$ and for multi-item auctions, $x_i \in [0, 1]^m$. A mechanism is called Bayesian Incentive Compatible (BIC) if it is every player's optimal strategy to report her true value, given all other players report truthfully. A mechanism is called Dominant Strategy Incentive Compatible (DSIC) if it is every player's optimal strategy to report her true value no matter what values are reported by the other players. Given any instance $\mathcal{I}$, let $\mathrm{Rev}(\mathcal{M}(\mathcal{I}))$ be the expected revenue generated by $\mathcal{M}$ and $\mathrm{OPT}(\mathcal{I})$ be the optimal BIC revenue of $\mathcal{I}$, i.e., the maximum expected revenue generated by BIC mechanisms. When $\mathcal{I}$ is clear from the context, we write OPT for short. A mechanism achieves c-approximation if for any instance $\mathcal{I}$,

$$\mathrm{Rev}(\mathcal{M}(\mathcal{I})) \geq \frac{\mathrm{OPT}(\mathcal{I})}{c}.$$

2.2 Query Complexity

In this work, we only allow the seller to access the prior distributions via two types of oracle queries: *value queries* and *quantile queries*. Given a distribution D over real numbers, in a value query, the seller sends a value $v \in \mathbb{R}$ and the oracle returns the corresponding quantile $q(v) \triangleq \Pr_{x\sim D}[x \geq v]$. In a quantile query, the seller sends a quantile $q \in [0, 1]$ and the oracle returns the corresponding value $v(q)$ such that $\Pr_{x\sim D}[x \geq v(q)] = q$. With *non-adaptive* queries, the seller first sends all his queries to the oracle, gets the answers back, and then runs the auction. The *query complexity* is the number of queries made by the seller.

Note that the answer to a value query is unique. The quantile queries are a bit tricky, as for discrete distributions there may be multiple values corresponding to the same quantile q, or there may be none. When there are multiple values, to resolve the ambiguity, let the output of the oracle be the largest one: that is, $v(q) = \arg\sup_z\{\Pr_{x\sim D}[x \geq z] \geq q\}$.[1] Note that for any discrete distribution D and quantile query $q > 0$, $v(q)$ is always in the support of D. Moreover, when $q = 0$, $v(q)$ may be $+\infty$.

Noisy Queries. In this paper, we also consider the model where the answers to the queries contain errors, i.e., noisy value queries and noisy quantile queries.[2]

Definition 1. *For distribution D and $\eta > 0$, a value query has η-noise if for any value $v \in \mathbb{R}$, the returned quantile is*

$$q \in \left[\frac{1}{1+\eta} \cdot \Pr_{x\sim D}[x \geq v], (1+\eta) \cdot \Pr_{x\sim D}[x \geq v]\right].$$

Similarly, a quantile query has η-noise if for any quantile $q \in [0, 1]$, the returned value v satisfies

$$v \in \left[\frac{1}{1+\eta} \cdot v', (1+\eta) \cdot v'\right],$$

and

$$v' = \arg\max_z\{\Pr_{x\sim D}[x \geq z] \geq q\}.$$

3 Main Results

We would like to understand both lower- and upper-bounds for the query complexity of approximately optimal Bayesian auctions. In this work, we will first consider single-item auctions and then extend our results to multi-item settings when the players' valuations are either unit-demand or additive. When the distributions are bounded within $[1, H]$ and the queries are precise, our results are summarized in Table 1. We also show that our query complexity results extend for arbitrary unbounded distributions that satisfy *small-tail assumptions*. A Bayesian auction instance $\mathcal{I}$ satisfies the small-tail assumption if there exists a function[3] $h : (0, 1) \to (0, 1)$ such that, for any constant $\delta_1 \in (0, 1)$ and any BIC mechanism $\mathcal{M}$, letting $\epsilon_1 = h(\delta_1)$, we have[4]

$$\mathop{\mathbb{E}}_{v\sim D} \mathbf{I}_{\exists i,j,q_{ij}(v_{ij})\leq\epsilon_1} \mathrm{Rev}(\mathcal{M}(v; \mathcal{I})) \leq \delta_1 \mathrm{OPT}(\mathcal{I}). \quad (1)$$

[1] The tie breaking rule here is chosen to simplify the exposition. All of our results extend to arbitrary tie breaking rules.

[2] In Definition 1, we define the noise as multiplicative errors. This choice is made since we focus on multiplicative approximations in our paper. Similar results can be obtained for additive error, which will not be elaborated in this paper.

[3] If computation complexity is a concern, then one can further require that the function is efficiently computable.

[4] In fact in our results for the single item auction, we only need a weaker condition where the inequality only need to hold for the Bayesian optimal mechanism OPT.

<table>
<tr><td rowspan="2"></td><td rowspan="2">Query Complexity</td><td colspan="3">Distributions</td></tr>
<tr><td colspan="2">Bounded in $[1, H]$</td><td>Unbounded & Small Tail</td></tr>
<tr><td rowspan="3">Auctions</td><td>Single-Item</td><td colspan="2">$\Theta(n\epsilon^{-1} \log H)$</td><td>$O(-n\epsilon^{-1} \log h(\frac{2\epsilon}{3(1+\epsilon)})$</td></tr>
<tr><td>Unit-Demand</td><td>$\forall c > 1$: $\Omega(\frac{mn \log H}{\log c})$</td><td>$\forall c > 24$: $O(\frac{mn \log H}{\log(c/24)})$</td><td>$\forall c > 24$: $O(-\frac{mn \log h(\frac{2c-48}{3c})}{\log(c/24)})$</td></tr>
<tr><td>Additive</td><td>$\forall c > 1$: $\Omega(\frac{mn \log H}{\log c})$</td><td>$\forall c > 8$: $O(\frac{mn \log H}{\log(c/8)})$</td><td>$\forall c > 8$: $O(-\frac{m^2 n \log h(\frac{c-8}{10c})}{\log(c/8)})$</td></tr>
<tr><td></td><td>Single-Item</td><td colspan="3">Regular Distributions: $\Omega(n\epsilon^{-1})$, $O(n\epsilon^{-1} \log \frac{n}{\epsilon})$</td></tr>
</table>

Table 1: Main results. Here $h(\cdot) < 1$ is the tail function in the small-tail assumptions. For single-item auctions, the revenue is a $(1 + \epsilon)$-approximation to the optimal BIC revenue, with ϵ sufficiently small. For multi-item auctions with unit-demand or additive valuations, the revenue is a c-approximation for some constant c. The upper bounds in the table hold for queries with small noise in the response while the lower bounds holds even for queries without noise. Note that the result of regular distributions doesn't require the distribution to be bounded.

Here $q_{ij}(v_{ij})$ is the quantile of v_{ij} under distribution D_{ij}, $\mathrm{Rev}(\mathcal{M}(v; \mathcal{I}))$ is the revenue of $\mathcal{M}$ under the Bayesian instance $\mathcal{I}$ when the true valuation profile is v, and $\mathbf{I}$ is the indicator function. For discrete distributions, Equation 1 is imposed on the ϵ_1 probability mass over the highest values. Intuitively, the small tail assumption assumes that the revenue contribution of any mechanism from high values, i.e., values with quantile smaller than ϵ, is sufficiently small. Similar small-tail assumptions are widely adopted in sampling mechanisms [Roughgarden and Schrijvers, 2016; Devanur *et al.*, 2016], to deal with irregular distributions with unbounded supports. Finally, we show that the revenue of the query mechanisms degrade proportionally to the amount of error contained in the queries to the valuation distributions.

Following the convention of the literature, the mechanism designer implements DSIC mechanisms while the benchmark is the optimal BIC revenue. It is known that the gap between the optimal DSIC mechanism and the optimal BIC mechanism for multi-item setting is bounded away from 1 [Yao, 2017], and the best known approximation ratio for DSIC mechanisms to the optimal BIC revenue in unit-demand and additive auctions are 24 and 8 respectively in the literature [Cai *et al.*, 2016]. The main focus of this paper is not the tightness of approximation ratios for those mechanisms, but the amount of information required for the mechanism designer to achieve such approximations.

Also note that our lower- and upper-bounds on query complexity are *tight* for bounded distributions. Actually, our lower-bounds allow the seller to make both value and quantile queries, and apply to any multi-player multi-item auctions where each player's valuation function is *succinct subadditive*. The lower-bounds also allow randomized adaptive queries and randomized mechanisms.

For the upper-bounds, all our query schemes are deterministic, non-adaptive, and only make one type of queries: value queries for bounded distributions and quantile queries for unbounded distributions with small-tail assumptions. We show that our schemes, despite of being very efficient, only loses a small fraction of revenue compared with when the seller has full access to the distributions.

4 More Related Work and Discussions

Sampling Mechanisms. A closely related area to our work is sampling mechanisms [Cole and Roughgarden, 2014; Devanur *et al.*, 2016; Babaioff *et al.*, 2018; Guo *et al.*, 2019; Hartline and Taggart, 2019; Allouah and Besbes, 2019]. It assumes that the seller does not know the distribution D but observes random samples from D before the auction begins. There are two branches of researches for sampling mechanisms. One branch focuses on designing mechanisms that minimize the approximation ratio given only one sample [Huang *et al.*, 2015; Allouah and Besbes, 2019] or two samples [Babaioff *et al.*, 2018; Daskalakis and Zampetakis, 2020]. The other one is referred as sample complexity, which measures how many samples the seller needs so as to obtain a good approximation to the optimal Bayesian revenue [Cai and Daskalakis, 2017; Guo *et al.*, 2019; Brustle *et al.*, 2020]. Most of the previous results focus on analyzing the sample complexity of DSIC mechanisms, which is similar to what we adopt in this paper. There are two exceptions. In [Gonczarowski and Nisan, 2017], the authors provide upper bounds on the sample complexity for $(1 + \epsilon)$-approximation to the optimal BIC revenue. However, the solution concept adopted there is ϵ-BIC, and the revenue guarantee under Bayesian Nash equilibrium for the mechanism proposed there is unclear. [Hartline and Taggart, 2019] resolve the issue for the single-item setting, and provide upper bounds on the sample complexity of non-truthful mechanisms for $(1 + \epsilon)$-approximation to the optimal BIC revenue, under the solution concept of Bayesian Nash equilibrium. The best-known sample complexity results are summarized in Table 2.

Oracle queries can be seen as *targeted samples*, where the seller actively asks the information he needs rather than passively learns about it from random samples. As such, it is intuitive that queries are more efficient than samples, but it is a priori unclear how efficient queries can be. Our results answer this question quantitatively and show that query complexity can be exponentially smaller than sample complexity: the former is *logarithmic* in the "size" of the distributions, while the latter is polynomial.[5] Finally, the design of query mechanisms facilitates the design of sampling mechanisms.

[5] For example, in the single-item auction, the sample complexity is $\tilde{O}(n^2 H \epsilon^{-2})$ when the valuations are bounded in $[1, H]$.

Auctions	Single-Item (regular)	Single-Item (bounded in $[1, H]$)	Unit-Demand (bounded in $[1, H]$)	Additive (bounded in $[1, H]$)
Sample Complexity	$\tilde{O}(n\epsilon^{-4})$	$\tilde{O}(nH\epsilon^{-3})$	$\tilde{O}(nm^2H^2\epsilon^{-2})$	$\tilde{O}(nm^2H^2\epsilon^{-2})$
Approximations	$1+\epsilon$	$1+\epsilon$	27	32

Table 2: Sample complexity in the literature. For multi-item auctions with unit-demand or additive valuation functions, the revenue has an extra ϵ additive loss. The results for single-item auctions are by [Devanur *et al.*, 2016], for unit-demand auctions are by [Morgenstern and Roughgarden, 2016], and for additive auctions are by [Cai and Daskalakis, 2017], respectively.

If the seller observes enough samples from D, then he can mimic quantile queries and apply query mechanisms.

Parametric Auctions. Parametric mechanisms [Azar and Micali, 2013; Carrasco *et al.*, 2018; Che, 2019] assume the seller only knows some specific parameters about the distributions, such as the mean, the median (or a single quantile), the variance or higher moments. Note that using quantile or value queries, one can get the exact value of the median and the approximate value of the mean, and then apply parametric mechanisms. Existing parametric mechanisms only consider single-parameter auctions. Since our mechanisms make non-adaptive oracle queries, our results imply parametric mechanisms in multi-parameter settings with general distributions, where the "parameters" are the oracle's answers to our query schemes. Our lower-bounds also imply that knowing only the median is not enough to achieve the same approximation ratios as we do.

Distributions within Bounded Distance. There are several papers in the literature that consider the model where the seller is given a distribution that is within a small distance to the true prior distribution. For single-item single-buyer settings, [Bergemann and Schlag, 2011] characterize the optimal robust monopoly pricing under the Prohorov distance. [Li *et al.*, 2019] consider the same problem with the earth-mover distance and extends the characterization to the multi-buyer setting. Recently, [Cai and Daskalakis, 2017] study the multi-item auctions under the Kolmogorov distance. The distinction between our model and those papers is that the learnt distributions from our query schemes may be far from the true prior in terms of those specified distances (e.g., Prohorov, Kolmogorov or earth-mover), and thus their mechanisms do not apply. On the other hand, although a distribution close to the true prior may be learnt via sufficiently many oracle queries given Prohorov distance, our results imply that the query complexity of this approach is not better than ours.

Menu Complexity. The complexity of auctions is an important topic in the literature, and other complexity measures such as menu complexity have been considered. Following the taxation principle [Hammond, 1979; Guesnerie, 1981; Hart and Nisan, 2013], defines the *menu complexity* of truthful auctions. For a single additive buyer, [Daskalakis *et al.*, 2013] show the optimal Bayesian auction for revenue can have an infinite menu size or a continuum of menu entries, and [Babaioff *et al.*, 2017] bound the menu complexity for approximating the optimal revenue. Recently, [Dobzinski, 2016] considers the taxation, communication, query and menu complexities of truthful combinatorial auctions, and shows important connections among them. The queries considered there are totally different from ours: we are concerned with the complexity of accessing the players' value distributions in Bayesian settings, while [Dobzinski, 2016] is concerned with the complexity of accessing the players' valuation functions in non-Bayesian settings.

Acknowledgements

We thank the editor and several anonymous reviewers for their helpful comments. Jing Chen is supported by NSF CAREER Award (No. 1553385). Bo Li is supported by RGC of HKSAR (No. PolyU 25211321), NSFC (No. 62102333) and The Hong Kong Polytechnic University (No. P0034420). Pinyan Lu is supported by Science and Technology Innovation 2030 –"New Generation of Artificial Intelligence" Major Project No.(2018AAA0100903), NSFC grant 61922052 and 61932002, Innovation Program of Shanghai Municipal Education Commission, Program for Innovative Research Team of Shanghai University of Finance and Economics, and the Fundamental Research Funds for the Central Universities. Part of this work was done when the first three authors were visiting Shanghai University of Finance and Economics.

References

[Allouah and Besbes, 2019] Amine Allouah and Omar Besbes. Sample-based optimal pricing. *SSRN 3334650*, 2019.

[Azar and Micali, 2013] Pablo Azar and Silvio Micali. Parametric digital auctions. In *4rd Innovations in Theoretical Computer Science Conference*, pages 231–232, 2013.

[Babaioff *et al.*, 2017] Moshe Babaioff, Yannai A Gonczarowski, and Noam Nisan. The menu-size complexity of revenue approximation. In *49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 869–877, 2017.

[Babaioff *et al.*, 2018] Moshe Babaioff, Yannai A Gonczarowski, Yishay Mansour, and Shay Moran. Are two (samples) really better than one? In *ACM Conference on Economics and Computation*, pages 175–175, 2018.

[Bergemann and Schlag, 2011] Dirk Bergemann and Karl Schlag. Robust monopoly pricing. *Journal of Economic Theory*, 146(6):2527–2543, 2011.

[Brustle *et al.*, 2020] Johannes Brustle, Yang Cai, and Constantinos Daskalakis. Multi-item mechanisms without item-independence: Learnability via robustness. In *21st*

ACM Conference on Economics and Computation, pages 715–761, 2020.

[Cai and Daskalakis, 2017] Yang Cai and Constantinos Daskalakis. Learning multi-item auctions with (or without) samples. In *58th Symposium on Foundations of Computer Science (FOCS'17)*, pages 516–527, 2017.

[Cai *et al.*, 2012] Yang Cai, Constantinos Daskalakis, and S Matthew Weinberg. Optimal multi-dimensional mechanism design: Reducing revenue to welfare maximization. In *53rd Symposium on Foundations of Computer Science (FOCS'12)*, pages 130–139, 2012.

[Cai *et al.*, 2016] Yang Cai, Nikhil R Devanur, and S Matthew Weinberg. A duality based unified approach to Bayesian mechanism design. In *48th Annual ACM Symposium on Theory of Computing (STOC'16)*, pages 926–939, 2016.

[Carrasco *et al.*, 2018] Vinicius Carrasco, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira. Optimal selling mechanisms under moment conditions. *Journal of Economic Theory*, 177:245–279, 2018.

[Chawla *et al.*, 2010] Shuchi Chawla, Jason D Hartline, David L Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *43th ACM Symposium on Theory of Computing (STOC'10)*, pages 311–320, 2010.

[Che, 2019] Ethan Che. Robust reserve pricing in auctions under mean constraints. *SSRN 3488222*, 2019.

[Cole and Roughgarden, 2014] Richard Cole and Tim Roughgarden. The sample complexity of revenue maximization. In *46th Annual ACM Symposium on Theory of Computing (STOC'14)*, pages 243–252, 2014.

[Cremer and McLean, 1988] Jacques Cremer and Richard P McLean. Full extraction of the surplus in Bayesian and dominant strategy auctions. *Econometrica*, 56(6):1247–1257, 1988.

[Daskalakis and Zampetakis, 2020] Constantinos Daskalakis and Manolis Zampetakis. More revenue from two samples via factor revealing sdps. In *21st ACM Conference on Economics and Computation*, pages 257–272, 2020.

[Daskalakis *et al.*, 2013] Constantinos Daskalakis, Alan Deckelbaum, and Christos Tzamos. Mechanism design via optimal transport. In *14th Conference on Electronic Commerce (EC'13)*, pages 269–286, 2013.

[Devanur *et al.*, 2016] Nikhil R. Devanur, Zhiyi Huang, and Christos-Alexandros Psomas. The sample complexity of auctions with side information. In *48th Annual ACM Symposium on Theory of Computing (STOC'16)*, pages 426–439, 2016.

[Dobzinski, 2016] Shahar Dobzinski. Computational efficiency requires simple taxation. In *57th Symposium on Foundations of Computer Science (FOCS'16)*, pages 209–218, 2016.

[Dwork, 2008] Cynthia Dwork. Differential privacy: A survey of results. In *International conference on theory and applications of models of computation*, pages 1–19. Springer, 2008.

[Gonczarowski and Nisan, 2017] Yannai A Gonczarowski and Noam Nisan. Efficient empirical revenue maximization in single-parameter auction environments. In *49th Symposium on Theory of Computing (STOC'17)*, pages 856–868, 2017.

[Guesnerie, 1981] Roger Guesnerie. On taxation and incentives: further remarks on the limits to redistribution. *University of Bonn*, 89, 1981.

[Guo *et al.*, 2019] Chenghao Guo, Zhiyi Huang, and Xinzhi Zhang. Settling the sample complexity of single-parameter revenue maximization. In *51st Annual ACM SIGACT Symposium on Theory of Computing*, pages 662–673, 2019.

[Hammond, 1979] Peter J Hammond. Straightforward individual incentive compatibility in large economies. *The Review of Economic Studies*, 46(2):263–282, 1979.

[Hart and Nisan, 2013] Sergiu Hart and Noam Nisan. The menu-size complexity of auctions. In *14th ACM Conference on Electronic Commerce*, pages 565–566, 2013.

[Hartline and Taggart, 2019] Jason Hartline and Samuel Taggart. Sample complexity for non-truthful mechanisms. In *20th ACM Conference on Economics and Computation*, pages 399–416, 2019.

[Huang *et al.*, 2015] Zhiyi Huang, Yishay Mansour, and Tim Roughgarden. Making the most of your samples. In *16th ACM Conference on Economics and Computation (EC'15)*, pages 45–60, 2015.

[Li *et al.*, 2019] Yingkai Li, Pinyan Lu, and Haoran Ye. Revenue maximization with imprecise distribution. In *18th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1582–1590, 2019.

[Morgenstern and Roughgarden, 2016] Jamie Morgenstern and Tim Roughgarden. Learning simple auctions. In *29th Conference on Learning Theory (COLT'16)*, pages 1298–1318, 2016.

[Myerson, 1981] Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.

[Ronen, 2001] Amir Ronen. On approximating optimal auctions. In *3rd ACM Conference on Electronic Commerce (EC'01)*, pages 11–17, 2001.

[Roughgarden and Schrijvers, 2016] Tim Roughgarden and Okke Schrijvers. Ironing in the dark. In *17th ACM Conference on Economics and Computation (EC'16)*, pages 1–18, 2016.

[Yao, 2017] Andrew Chi-Chih Yao. Dominant-strategy versus bayesian multi-item auctions: Maximum revenue determination and comparison. In *18th Conference on Economics and Computation (EC'17)*, pages 3–20, 2017.